

Asset Allocation in Finance: A Bayesian Perspective

Eric Jacquier

HEC Montreal and MIT Sloan

Nicholas G Polson

Booth School of Business

*University of Chicago**

First Draft: March 2011

This Draft: August 2011

Abstract

We survey asset allocation in finance from a Bayesian decision-theoretic perspective. We study an investor who wishes to maximize the expected long-run growth of the market. With the aid of Stein's lemma we derive the Kelly criteria for optimal bet size and Merton's allocation rule for risky stocks. We therefore provide an equivalence between these two criteria. Bayesian inference naturally determines the inputs for this analysis, namely, the expected excess return and volatility of the risky asset. Extensions to exchangeable returns where the investor learns about the probability of success illustrate the feature that risk-averse investors are still willing to hold a small proportion of a risky asset even though the odds are unfavorable at the current time. The option value of future learning leads one to a positive allocation. We conclude with directions for future research.

*Polson is Professor of Econometrics and Statistics at the Chicago Booth School of Business. email: ngp@chicagobooth.edu. Jacquier is associate professor of finance, visiting at MIT Sloan on leave from HEC Montreal, email: jacquier@mit.edu. Jacquier acknowledges acknowledges support from the HEC Montreal professorship in derivative securities.

1 Introduction

Bayesian methods have long played a role in finance and asset allocation since the seminal work of de Finetti (1941) and Markowitz (2006). In this paper, we show how the principle of maximum expected utility (MEU) (Ramsey, 1926, Savage, 1956, Bernardo and Smith, 2000) together with Stein’s lemma for stochastic volatility distributions (Gron, Jorgensen and Polson, 2011) solves for the optimal asset allocation. Stein’s lemma provides the solution to the first order condition that accompanies MEU. The optimal asset allocation problem couched in equilibrium then leads to models such as the Capital Asset Pricing Model (CAPM) or Merton’s inter-temporal asset pricing model (ICAPM).

We consider an investor who wishes to invest in the risky asset in order to maximize the expected utility of her resulting wealth. Under logarithmic utility, this leads to the famous Kelly criterion which maximizes the expected long-growth rate of the market. We review the link between the Kelly rule (Kelly, 1956) and the Merton optimal asset allocation (Merton, 1969). We illustrate their implementation for a discrete binary setting as well as a historical stock returns on the S&P500. Under a CRRA utility with Stein’s lemma we show how to derive fractional versions of these allocation rules where the amount allocated is normalized by the investor’s relative risk aversion.

The rest of the paper is as follows. The next section reviews the impact of Bayesian thinking in models of finance: asset pricing equilibrium models; how agents learn from prices; properties of returns data including stochastic volatility. Section 2 views asset allocation from a Bayesian decision theoretic perspective (Bernardo and Smith, 2000). Section 3 studies the problem of maximizing the expected long-run growth rate and derive the classic Kelly and Merton allocation rules. Section 3 describes methods for estimating this long-run growth rate. Section 4 describes estimation methods for long-run asset allocation. Section 5 considers extensions to Bayesian dynamic learning (Bellman and Kalaba, 1956) and time-varying investment opportunity sets (Ferguson and Gilstein, 1985). When investors are faced with a return distribution that is an exchangeable Beta-Binomial process, the effect of dynamic learning leads makes them willing to invest a small amount of capital to current returns that have a negative expectation even though they are averse to risk. This is due to the fact that they might learn that the investment opportunity set improves in the future and this is taken account of in the Bayesian MEU solution. Finally, Section 6 concludes.

1.1 Bayesian methods in Finance

Aside from the traditional estimation methods, Bayesian thinking underpins finance in a number of empirical and theoretical ways.¹ For example, many of the theoretical de-

¹Bayesian MCMC and particle methods are prevalent in empirical finance, see e.g. Jacquier, Polson and Rossi (1994, 2004). Johannes and Polson (2010) discuss MCMC methods in financial econometrics, Brandt (2009) describes portfolio choice problems and Polson and Tew (2000) provide an empirical analysis of the S&P500 stock index. Jacquier and Polson (2011) provide a recent survey of Bayesian methods in finance.

velopments in finance rely on Bayesian learning by agents. Here we discuss applications to learning from prices; understanding bubbles and speculative behaviour and finally, measuring the stochastic volatility of asset returns (Black 1976), which has been an empirical feature of markets for many centuries (see, for example, Harrison, 1998 and 2001).

Learning from Prices Bayesian methods are designed for learning, and market equilibrium occurs after individuals with differences in opinion have an incentive to trade and finally agree on a price. The differences in opinion literature has a long history dating to DeGroot (1973), Geanakoplos and Polemanokos (1982), Geanakoplos and Sebenius (1989), Harris and Raviv (1993), Hong and Stein (2005) and Varian (1985, 1987). Odean (1998) and Barberis and Thaler (2003) discuss investor behavior in financial markets.

A basic argument is as follows. Let y denote an observed signal. The insight is that observing prices $P(y)$ will change a trader's beliefs. One caveat is that she has to be able to coherently update her probabilities via Bayes theorem. Grossman (1976, 1981) summarizes the logic clearly as follows:

A smart trader t might even say to himself: let all the other traders naively use their own information. I will wait until the market clears, and after observing the current realization of $P(y)$, make my purchases of commodities to maximize $E[U(W)|y_t, P(y)]$. Since I am a price taker, I will expect to do better than by trading now and maximizing $E[U(W)|y_t]$

Market efficiency arguments follows a similar route, see Grossman and Stiglitz (1980). Let $\pi(y)$ be the vector of probabilities that the trader holds for the states. If the observed prices $P(\pi(y))$ are invertible in π , the trader can back out everybody else's beliefs. Then implementing the principle of maximum expected utility

$$E[U_t(W_t)|y_t, P(\pi(y))] \equiv E[U_t(W_t)|y]$$

The trader who only has information y_t can still act *as if* he had full information set of all investors signals y . Hence, in an efficient market prices are fully revealing of trader's beliefs.

Bubbles and Speculative Behaviour Many models have been proposed to update bubbles and speculative behaviour, see, for example, the seminal work of Harrison and Kreps (1978) and Scheinkman and Xiong (2003). Stein (2009) provides a recent review. Bubbles are of fundamental interest. Early models in this literature "explained" bubbles by incorporating agents who are persistently "naive" and do not update beliefs coherently. The recent literature has focused more on the dynamics of Bayesian learning and belief structures. For example, Morris (1996) allows agents to do some beta-binomial updating of their beliefs, and Heaton and Polson (2011) use prices to learn whether one is smart or dumb money in a market.

One of the most famous bubbles was the South Sea bubble of 1720. The English parliament passed the Bubble Act in the following year 1721 trying to ban future bubbles. At the height of the bubble, the price-earnings multiple of the stock was between 150 and 190 based on current earnings. In the first six month of 1720 the South sea stock rose 500 percent. The major "news" being only the large demand for their shares. A number of investors lost large sums of money. Sir Isaac Newton (1642-1729) who was Master of the Mint from 1699 to 1729 and who put England on the Gold standard in 1717, reportedly lost £20,000 in the Bubble. He was quoted as saying *I can predict the motion of Heavenly bodies but not the behavior of the stock-market.*

Return Dynamics and Stochastic Volatility One of the interesting features of stock return series is stochastic volatility. Volatility is time-varying, persistent, mean-reverting and in some instances has jumps (Jacquier, Polson, Rossi, 1994, 2005, Eraker, Johannes and Polson, 2003). Harrison (1998, 2005) shows that the eighteenth and twentieth century time series distributions of stock returns are statistically very similar, and both exhibit stochastic volatility. In particular, he analyzes biweekly prices and returns for shares of the London stock market index and the Dutch East India company from 1723-1794.

We now turn to the problem of asset allocation.

2 Asset Allocation

The asset allocation problem can be described as a Bayesian decision-theoretic problem (Bernardo and Smith, 2000) as follows. An investor has initial wealth W_0 that can be invested in a risk-free bond or a risky stock. The end of period wealth is $W = W_0 \{(1 - \omega)r_f + \omega R\}$ where r_f is the return on the risk-free rate and R is the return on the risky asset. The Bayesian investor maximises expected utility, namely $\max_{\omega} \mathbb{E}[U(W(\omega))]$ which leads to a first order condition of the form:

$$E[U'(W)(R - r_f)] = 0$$

Applying the definition of covariance yields

$$Cov[U'(W), R - r_f] + E[U'(W)] E[R - r_f] = 0$$

$$\omega E[U''(W)] var(R) + E[U'(W)] E[R - r_f] = 0.$$

where we have used Stein's lemma for a differentiable function $g(X)$ where $\mathbb{E}|g'(X)| < \infty$ then $cov(g(X), X) = \mathbb{E}(g'(X)) var(X)$. This has a solution for the optimal weight

$$\omega^* = \frac{1}{\gamma} \left(\frac{\mathbb{E}[R] - r_f}{Var[R]} \right)$$

where $\gamma = -\mathbb{E}[U'']/\mathbb{E}[U']$ is the agents risk aversion parameter. This is a form of the famous Merton optimal asset allocation result.

This result can be extended to allow for stochastic volatility in the risky asset return. Let $(X|V) \sim \mathcal{N}(0, V)$ where $V \sim p(V)$. The equivalent Stein result is

$$Cov[g(X), Y] = \mathbb{E}^Q[g'(X)] Cov[X, Y]$$

Here \mathbb{E}^Q is taken with respect to the distribution $q(V) = Vp(V)/\mathbb{E}[V]$. This is size-biasing of the original volatility distribution $p(V)$. If $Y = X$, we have $Cov[g(X), X] = \mathbb{E}^Q[g'(X)] var[X]$. The optimal allocation rule is then

$$\omega_{SV}^* = \frac{1}{\Gamma^Q} \left(\frac{\mathbb{E}[R] - r_f}{Var[R]} \right)$$

where $\Gamma^Q = -\mathbb{E}^Q[U'']/\mathbb{E}[U']$ is the volatility adjusted risk aversion. See Gron et al (2011) for an analysis of comparative statics.

In this one period setting, given that agents adopt the optimal allocation rule, one can then show that an equilibrium Capital Asset Pricing Model (CAPM) holds. Specifically, under stochastic volatility (SV), the expected return $\mathbb{E}(R_j)$ on the j th security is given by

$$\mathbb{E}(R_j) - r_f = \beta_j (\mathbb{E}(R_m) - r_f)$$

where r_f is the risk-free rate and $\beta_j = cov(R_j, R_m)/var(R_m)$ is the risk premium.

A commonly used class of utility function is the constant relative aversion (CRRA) class. It has the advantage that the optimal rule is unaffected by wealth effects. It takes the form

$$U_\gamma(W) = (W^{1-\gamma} - 1)/(1 - \gamma)$$

The special case $U(W) = \log(W)$ for $\gamma = 1$ plays a central role in growth rate analysis. It leads to maximising the expected long-run rate of growth, namely

$$\max_{\omega} \mathbb{E}(\log W_T | W_0 = x)$$

We will now solve for this rule and derive the Kelly criterion and Merton's rule.

3 Maximising Expected Long-Run Growth

3.1 Kelly Rule and Merton's Optimal Allocation

The Kelly Criterion corresponds to the following Bayesian decision problem under binary uncertainty. Consider a sequence of i.i.d. bets where

$$p(X_t = 1) = p \text{ and } p(X_t = -1) = q = 1 - p$$

The investor who maximises $\mathbb{E}(\log W_T | W_0 = x)$ uses a myopic-rule with weight

$$\omega^* = p - q = 2p - 1$$

Indeed, we can see that maximising the expected long-run growth rate leads to the solution

$$\begin{aligned} \max_{\omega} \mathbb{E}(\ln(1 + \omega W_T)) &= p \ln(1 + \omega) + (1 - p) \ln(1 - \omega) \\ &\leq p \ln p + q \ln q + \ln 2 \text{ and } \omega^* = p - q \end{aligned}$$

If you believe the event is certain i.e. $p = 1$, then you bet all of your wealth and *a priori* you think you are certain to double your wealth. On the other hand, if you think the bet is fair, i.e. $p = \frac{1}{2}$, you bet nothing, $\omega^* = 0$, as you are risk-averse.

We will use the following notation. Let p denote the probability of a gain and $O = (1-p)/p$ the odds. We can generalize the rule to the case of asymmetric payouts (a, b) where

$$p(X_t = 1) = p \text{ and } p(X_t = -1) = q = 1 - p$$

Then the objective expected utility function is

$$p \ln(1 + b\omega) + (1 - p) \ln(1 - a\omega)$$

with optimal solution

$$\omega^* = \frac{bp - aq}{ab} = \frac{p - q}{\sigma}$$

If $a = b = 1$ this reduces to the pure Kelly criterion.

A common case occurs when $a = 1$. We can now interpret b as the odds O that the market is willing to offer the invest if the event occurs and so we write $b = O$. The rule becomes

$$\omega^* = \frac{p \cdot O - q}{O}$$

We now provide a counter-intuitive example.

Market	You	p	ω^*
4/1	3/1	1/4	1/16
12/1	9/1	1/10	1/40

Table 1: Kelly rule

Example: Betting Consider the following two betting situations described in Table 1. Suppose that there are two possible market opportunities: one where it offers you 4/1 when you have personal odds of 3/1 and, second, when it offers you 12/1 and you think the odds are 9/1. In expected return these two scenarios are identical both offering a 33% gain. In terms of maximizing long-run growth, however, they are not identical. From Table 1, the Kelly criteria advises an allocation that is twice as much capital to the lower odds proposition: 1/16 weight versus 1/40.

Specifically, we have the following optimal weight calculation $\omega^* = (pO - q)/O$ with allocations

$$\frac{(1/4) \times 4 - (3/4)}{4} = \frac{1}{16} \text{ and } \frac{(1/10) \times 12 - (9/10)}{12} = \frac{1}{40}$$

respectively. We now turn to the continuous-time setting and find Merton's rule.

Continuously Compounded Returns: In a continuous-time setting, let μ denote the (excess) expected return, σ the volatility and γ the risk aversion. Suppose that the risky asset follows a Black-Scholes geometric Brownian motion model

$$dS_t = S_t (\mu dt + \sigma dB_t)$$

for a constant volatility σ . Then the value of the asset at time T is

$$S_T = S_0 \exp \left\{ \left(\mu - \frac{1}{2} \sigma^2 \right) T + \sigma \sqrt{t} Z \right\}$$

where $Z \sim N(0, 1)$. This model implies that returns are log-normally distributed.

Now consider the evolution of wealth for an investor who keep a constantly rebalanced weight ω allocated to the risky asset. Her wealth W_t^ω now evolves according to $dW_t^\omega/W_t^\omega = \omega \mu dt + \omega \sigma dB_t$. Here we write $\mu \leftarrow \mu - r_f$ as the expected excess return of the asset for notational convenience.

It can be shown, for example, by Ito's lemma, to have a value at T

$$W_T^\omega = W_0 \exp \left\{ \left(\omega \mu - \frac{1}{2} \omega^2 \sigma^2 \right) T + \omega \sigma \sqrt{t} Z \right\}$$

Long-run Growth rate This leads to a natural definition of the growth rate as

$$G(\omega) \triangleq \omega\mu - \frac{1}{2}\omega^2\sigma^2$$

Maximize the growth rate with respect to ω leads to the optimal allocation

$$\omega^* = \frac{\mu}{\sigma^2}$$

This is known as the pure Kelly rule. Fractional Kelly rules are, as their name indicates, rules that allocate a fraction of the Kelly rule to the risky asset. The maximization of expected utility when the investor has a constant relative risk aversion γ , and returns are continuously compounded leads to Merton's optimal asset allocation:

$$\mathbb{E}[U(W_T^\omega)] = W_0^{1-\gamma} \exp \left\{ \left(\mu\omega\gamma - \frac{1}{2}\omega^2\gamma^2\sigma^2 \right) T \right\} \text{ and } \omega^* = \frac{1}{\gamma} \frac{\mu}{\sigma^2}$$

This can be seen as a utility interpretation of fractional Kelly rules. This analysis also agrees with Stein's lemma approach in the previous section.

We can analytically compute the growth rate of wealth at the optimal allocation as:

$$G(\omega^*) = \frac{1}{\gamma} \left(1 - \frac{1}{2\gamma} \right) \left(\frac{\mu}{\sigma} \right)^2$$

One sees that if $\gamma < \frac{1}{2}$, a case of agents with extremely low risk aversion, then $G(\omega^*) < 0$. Moreover, we can see how estimation risk affects the growth rate: if estimation error mistakenly lets leads to using twice the optimal rule, then we have lost all of the growth of our portfolio as $G(2\omega^*) = 0$.

We now provide an equivalence with the Kelly criterion by scaling wealth in terms of volatility units. In the continuous case, let the investor be faced with a sequence of returns where the payout has expectation $E(R_T) = \mu$

$$p(R_T = \sigma) = p = \frac{1}{2} + \frac{1}{2} \frac{\mu}{\sigma} \text{ and } p(R_T = -\sigma) = q = \frac{1}{2} - \frac{1}{2} \frac{\mu}{\sigma}$$

Given (μ, σ) , the optimal Kelly rule is

$$\omega^* = \frac{p - q}{\sigma} = \frac{\mu}{\sigma} \cdot \frac{1}{\sigma} = \frac{\mu}{\sigma^2}$$

This provides our equivalence with Merton's rule.

Example S&P500: Consider a simple example of logarithmic utility (CRRA with $\gamma = 1$). This is a pure Kelly rule. We assume iid log-normal stock returns with an annualized expected

excess return of 5.7% and a volatility of 16% which is consistent with long-run equity returns. In our continuous time formulation $\omega^* = 0.057/0.16^2 = 2.22$ and the Kelly criterion which implies that the investor borrows 122% of wealth to invest a total of 220% in stocks. This is the risk-profile of the Kelly criterion. One also sees that the allocation is highly sensitive to estimation error in $\hat{\mu}$. We consider dynamic learning in a later section and show how the long horizon and learning affects the allocation today.

The fractional Kelly rule leads to a more realistic allocation. Suppose that $\gamma = 3$. Then the informational ratio is

$$\frac{\mu}{\sigma} = \frac{0.057}{0.16} = 0.357 \text{ and } \omega^* = \frac{1}{3} \frac{0.057}{0.16^2} = 74.2\%$$

An investor with such a level of risk aversion then has a more reasonable 74.2% allocation.

This analysis ignores the equilibrium implications. If every investor acted this way, then this would drive up prices and drive down the equity premium of 5.7%.

4 Long Run Asset Allocation

Discussions of long term investment policy, e.g., retirement, revolve around measures of the expected long-term return on a risky portfolio such as the global market index. In this section we review how optimal Bayesian asset allocation leads to a estimate of long run expected returns with very attractive properties. We first give a background on the estimates resulting from the classical literature, and their potential shortcoming.

4.1 Background on Classical Estimation

To concentrate on the main issue, that of the uncertainty in the mean, we assume that variance is known and that returns are not auto-correlated, see Jacquier et al. (2005) and Jacquier (2008) for a robustness analysis. Recall that, as financial returns exhibit very little auto-correlation, there is no gain in observing the data more frequently. In this context, we therefore consider a sample of T annual i.i.d. log-normal returns $\log(1 + R) \sim N(\mu, \sigma)$. The reader will quickly see that she can adapt the discussion to any estimator of the mean return given its associated variance. The long-term, H -period return, or wealth return per dollar invested is log-normal. Long-term investors and policy-makers seek an estimate of its expected compound return

$$E(V_H) = e^{H(\mu + 0.5\sigma^2)}. \tag{1}$$

The MLE estimate of μ is the sample mean, it corresponds to the Bayesian posterior mean with no prior information. For long-term forecasts, practitioners used to choose a point estimate by compounding the sample geometric return $G = \frac{1}{T} \log \frac{P_T}{P_1}$. This amounts to estimating $E(V_H)$ by $e^{\hat{\mu}H}$. Academics, however, tended to substitute $\hat{\mu}$ in the theoretical

expectation (1), often invoking a maximum likelihood (ML) justification, where the estimator of a function is approximated by the function of the estimator. This second suggestion is equivalent to the compounding H times of the arithmetic sample mean. The difference in these two estimates becomes very large in the long run. Using Siegel's (1994) geometric and arithmetic averages of 7% and 8.5%, the two approaches grow \$1 to \$160 versus \$454 over 75 years. Jacquier et al. (2003) show that the ML approach makes little sense; as H is often of a magnitude comparable to T , the ML estimator suffers from an enormous upward bias due to the Jensen effect. An unbiased estimator is $U = \exp\{H(\hat{\mu} + 0.5\sigma^2(1 - \frac{H}{T}))\}$, a simple log-linear combination of the geometric and arithmetic estimators. Observe how the compounding factor is linearly decreasing in H . Jacquier et al. (2005) argue that there is little justification for mere unbiasedness, and derive a minimum mean squared error classical estimator of $E(V_H)$:

$$M = e^{H(\hat{\mu} + 0.5\sigma^2(1 - 3\frac{H}{T}))}. \quad (2)$$

These estimators nest the ML (aka arithmetic) estimator, only justified with extremely small $\frac{H}{T}$, and the geometric estimator, itself only justified when H happens to be equal to $\frac{T}{3}$. For both U and M , the compounding factor decreases linearly with the horizon H , with a downward penalty increasing with $\frac{H}{T}$. This is desirable because, (1) H affects the amount of compounding that magnifies (by Jensen effect) the upward bias due to the estimation error in μ , and (2) T reduces the uncertainty in μ . For realistic values of μ, σ, T, H , the penalty is quite severe; longer term investors must be far more modest in their predictions than shorter term investors, given the same base estimator of μ .

However, since the compounding factor decreases linearly, it can lead for to expected returns lower than the risk-free rate, even negative, for long enough horizons. This feature of the estimator makes no economic sense; any discount factor should tend to no lower than the risk-free rate as the horizon increases. We now turn to a simple Bayesian estimator based on optimal asset allocation, which naturally incorporates this desired feature.

4.2 Bayesian Estimation of the Long Run Expected Return

Parameter uncertainty has an important impact on the optimal allocation. This has been recognized since Brown (1976). Intuitively, the proper distribution for the investor to consider is the predictive density, which gets inflated by the uncertainty in the mean. Relative to the case with known parameters, the optimal allocation is then lower in the risky asset.

Consider first the classic Merton (1969) asset allocation, it is a framework with one risky asset with a i.i.d. normally distributed log-return $N(\mu, \sigma^2)$, a risk-free return r_f , and a power utility with risk aversion γ . If all parameters are known, Merton (1969) shows that, with continuous re-balancing, the horizon H is irrelevant and drops out of the computation.

He derives the now well-known optimal allocation,

$$w^* = \frac{\alpha - r_f}{\gamma\sigma^2} \quad (3)$$

where $\alpha = \mu + 0.5\sigma^2$ is the expected arithmetic one-period return. To obtain this result, assuming constant rebalancing to an allocation w , it is easy to show first that the random H period return is log normally distributed as

$$\log(V_H|\alpha, \sigma) \sim N \left[(r_0(1-w) + w\alpha - 0.5w^2\sigma^2)H, w^2\sigma^2H \right]. \quad (4)$$

The expected utility of this random wealth is then:

$$E[U(V_H)] = \frac{1}{1-\gamma} \exp \left[(1-\gamma)H(r_0 + w(\alpha - r_0) - 0.5w^2\sigma^2 + 0.5(1-\gamma)w^2\sigma^2) \right]. \quad (5)$$

The maximum likelihood estimator of μ is the standard average of past log-returns $\hat{\mu} \sim N(\mu, \sigma^2/T)$. With diffuse priors, the Bayesian posterior of μ is numerically equivalent, albeit with the well-known difference in interpretation. It is clear that simply substituting $\hat{\alpha}$ in eqn (5) or (3) is not the optimal solution. Jacquier (2008) derives the Bayesian optimal allocation with uncertainty in μ . The investor must consider the predictive utility, i.e., the utility of the predictive density of the H period return. Because the integrations over the parameter and over the distribution of returns can be exchanged, one can also view this as integrating μ out of the conditional expected utility in (5), using its posterior distribution. The expected (predictive) utility becomes:

$$E[U(V_H)] = \frac{1}{1-\gamma} \exp \left[(1-\gamma)H \left[r_0 + w(\hat{\alpha} - r_0) - 0.5w^2\sigma^2 + 0.5(1-\gamma)w^2\sigma^2 \left(1 + \frac{H}{T} \right) \right] \right]. \quad (6)$$

Note that α has been replaced by its posterior mean $\hat{\alpha}$, and there is a new term in H/T at the end. Maximizing the expected utility in (6), Jacquier (2008) finds the optimal asset allocation under uncertainty:

$$w^* = \frac{\hat{\alpha} - r_0}{\sigma^2 \left[\gamma \left(1 + \frac{H}{T} \right) - \frac{H}{T} \right]}. \quad (7)$$

Figure 1 plots the optimal allocation versus the horizon for an investors with a relative risk aversion of 4 and standard values of the mean and variance of the market return. It shows that the effect of uncertainty on the optimal allocation is important for realistic values of μ, r_f, σ, γ .

Insert Figure 1 here.

This optimal allocation in (7) implies an estimate of the expected long-term return under uncertainty. For this investor with risk aversion γ , denote α^* this *risk-adjusted* future

expected return, where the risk is estimation risk. For this risk-adjusted return, the optimal allocation would be the Merton allocation in (3). Equating the two allocations, we find that this risk-adjusted estimate is:

$$\alpha^* - r_f = \frac{\hat{\alpha} - r_f}{1 + \frac{H}{T}(1 - \frac{1}{\gamma})}. \quad (8)$$

It can be rewritten in terms of μ if desired. Figure 2 shows this estimate versus the horizon. Even with a moderate risk aversion of 4, it reinforces the stark (low expected return) prediction of the minimum mean squared error estimate in (2).

Insert Figure 2 here

This estimate is optimal for a given investor with risk aversion γ . Now, not only investors with different horizons but also investors with different risk aversion, formulate different optimal point forecasts. Note that the estimate in (8) can never be below the risk-free rate. This is to be expected since it is consistent with the Bayesian optimal allocation which has a lower bound of zero for the very long run.

5 Dynamic Bayes Learning

When the opportunity set varies in time, the agent that maximising long run growth will take this into account in his decision rule today. The Bayesian dynamic portfolio problem dates to Bellman and Kalaba (1956, 1958). Here we will follow the analysis of Ferguson and Gilstein (1985) and consider the extensions to CRRA utility and to exchangeable return distributions where the agent will engage in dynamic learning.

Remember that in an i.i.d. discrete setting, the classic Kelly rule is for $p > \frac{1}{2}$,

$$\omega^* = p - q = 2p - 1$$

and zero otherwise. The investor never bets on outcomes which have unfavorable (given his views) odds, or even fair odds due to risk aversion.

Now if returns are still independent but have time-varying known probabilities p_t we can still solve for the optimal allocation: for $p_t > \frac{1}{2}$ invest

$$\omega_t^* = \frac{p_{t-1}^{\gamma-1} - q_{t-1}^{\gamma-1}}{p_{t-1}^{\gamma-1} + q_{t-1}^{\gamma-1}}$$

If p_t is unknown, a myopic Bayes rule is to track the sufficient statistics under a Beta prior

$Be(\alpha, \beta)$ distribution with T observations $y = (y_1, \dots, y_T)$ as

$$\mathbb{E}(p|y) = \frac{\alpha \sum_{t=1}^T y_t}{T + \alpha + \beta}$$

and uses a myopic plug-in rule by replacing p with $E(p|y^T)$.

To be fully Bayesian, we consider the exchangeable case where the investor can learn from experience. The solution of maximizing discounted expected utility has a very different solution. We need to solve for the value function using Bellman's equation and take account of the fact that *in the future* he will be using Bayes rule and learning from prices.

Bayes Learning Consider a conjugate Bernoulli-Beta model where $(y|p) \sim Ber(p)$ and prior distribution $(p|a, b) \sim Be(a, b)$ for given hyperparameters.

Value function If we consider the case of power utility, the agent solves

$$\max_{\omega} \mathbb{E} [U(W_T)|W_0 = x] = x^{1-\gamma} V_T(a, b)$$

where V_T is the value function at period n .

Let $f_T(x|\alpha, \beta)$ denote the optimal value of $\mathbb{E} (W_T^{1-\gamma}|W_0 = x)$ when the prior distribution is $p \sim Be(\alpha, \beta)$. The posterior is then $Be(\alpha + 1, \beta)$ given a success and $Be(\alpha, \beta + 1)$ given a failure. At the initial condition, we have $f_0(x|\alpha, \beta) = x^{1-\gamma}$. We can then solve for the value function recursively using the identity

$$\begin{aligned} f_1(x|\alpha, \beta) &= \max_{0 \leq b \leq x} \left\{ \frac{\alpha}{\alpha + \beta} (x + b)^{1-\gamma} + \frac{\beta}{\alpha + \beta} (x - b)^{1-\gamma} \right\} \\ &= \max_{0 \leq c \leq 1} \left\{ \frac{\alpha}{\alpha + \beta} (1 + c)^{1-\gamma} + \frac{\beta}{\alpha + \beta} (1 - c)^{1-\gamma} \right\} \\ &= x^{1-\gamma} V_1(\alpha, \beta) \end{aligned}$$

Following recursively, we find

$$V_k(a, b) = \max_{0 \leq c \leq 1} \left\{ \frac{a}{a + b} V_{k-1}(a + 1, b)(1 + c)^{1-\gamma} + \frac{b}{a + b} V_{k-1}(a, b + 1)(1 - c)^{1-\gamma} \right\}$$

The optimal allocation is then

$$\omega_k(a, b) = \left(\frac{1 - w}{1 + w} \right)^+ \quad \text{where } w = \left(\frac{b V_{k-1}(a, b + 1)}{a V_{k-1}(a + 1, b)} \right)^{\gamma^{-1}}, \gamma > 1$$

In the myopic Kelly case where $\gamma = 1$, we have the rule $\omega_k(\alpha, \beta) = 0$ for $\alpha < \beta$ as one might expect. You will not invest if your prior mean says the bet is unfair.

This is not true in general though due to the learning effect. For example, with a uniform prior

$$\begin{aligned}w_{10}(1, 1) &= 0.875 \text{ and } w_9(1, 2) = 0 \\w_{10}(1, 1) &= 0.973 \text{ and } w_9(1, 2) = 0.117\end{aligned}$$

and the investor allocates a large portion to the risky asset. Even in the case where he sees a failure and he thinks that the odds are two to one against him, he is still willing to invest $w_9(1, 2) = 0.117$ that is 11.7% of his wealth.

We therefore have the important effect of Bayesian learning on asset allocation, that the agent will invest in the first period even though the expected return is negative and s/he is risk averse!

6 Discussion

Bayesian thinking is central to finance and asset allocation. Stein's lemma provides a useful tool for analyze first order maximum expected utility conditions. The classic Kelly and Merton allocation rules correspond to Bayes rules where the investor maximizes the expected long-run growth rate of accumulation. These rules are sensitive to estimation risk and Bayesian estimation methods provide the necessary inputs for expected returns and volatility.

Under the maximum growth condition when returns are independent these rules are myopic. Investors will not be willing to allocate any capital to unfavorable bets. However, when the invest opportunity set is extended to allow for Bayesian learning with exchangeable return distributions, we see that risk-averse investors are willing to hold a small amount of the risky asset in the hope that they will learn in the future that conditions are improved in the future.

There are many avenues for future research. One currently active area is to incorporate predictability into the return distribution. The dividend yield regressor (Barberis, 2000) together with stochastic volatility (Johannes, Korteweg and Polson, 2011) plays a central role. In general the portfolio problem incorporates economically significant hedging demand and sensitivity to priors is an important issue.

7 References

- Barberis, N. (2000). Investing in the Long Run when Returns are Predictable. *Journal of Finance*, 55, 225-264.
- Barberis, N. and R. Thaler. (2003). A Survey of Behavioral Finance. *Handbook of the Economics of Finance* (eds Constantinides et al), 1053-1128.
- Bellman, R. and R. Kalaba (1956). On the role of dynamic programming in statistical communication theory. *IRE Transactions on Information Theory*, IT-3, 197-203.
- Bellman, R. and R. Kalaba (1958). On Communication processes involving Learning and Random Duration. *IRE National Convention Record*, 4, 16-20.
- Bernardo, J. and A.F.M. Smith (2000). *Bayesian Theory*. Wiley.
- Black, F. (1976). Studies of Stock Market Volatility Changes. *Proceedings of Journal of American Statistical Association*, 177-181.
- Borel, E. (1924). Apropos of a Treatise on Probability. *Revue Philosophique*. Reprinted in *Studies in Subjective Probability*, Kyburg and Smokler (1980) (eds).
- Brandt, M. (2009). Portfolio Choice Problems. *Handbook of Financial Econometrics* (eds Ait-Sahalia and L. P. Hansen), 269-336.
- Brown, S. (1976). Optimal Portfolio Choice Under Uncertainty: A Bayesian Approach. *Ph.D, University of Chicago*.
- de Finetti, B. (1941). Il problema dei Pieni. Reprinted: *Journal of Investment Management*, 4(3), 19-43.
- DeGroot, M.H. (1973). Reaching a consensus. *Journal of the American Statistical Association*, 69, 118-121.
- Eraker, B., M. Johannes and N.G. Polson (2003). The Impact of Jumps in Volatility and Returns. *Journal of Finance*, 58(3), 1269-1300.
- Ferguson, T.S. and C.Z. Gilstein (1985). A General Investment Model. *Technical Report, UCLA*.
- Geanakoplos, J. and H. Polemarchakis (1982). We can't disagree forever. *Journal of Economic Theory*, 192-200.
- Geanakoplos, J. and J. Sebenius (1983). Don't bet on it: a note on contingent agreements with asymmetric information. *Journal of the American Statistical Association*, 78, 224-226.

- Gron, A., B. Jorgensen and N.G. Polson (2011). Optimal Portfolio Choice and Stochastic Volatility. *Applied Stochastic Models*.
- Grossman, S. (1976). On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information. *Journal of Finance*, 31, 573-85.
- Grossman, S. (1978). Further Results on the Informational Efficiency of Competitive Stock Markets. *Journal of Economic Theory*, 18, 81-101.
- Grossman, S. J. and J.E. Stiglitz. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, 70, 393-408.
- Harrison, P. (1998). Similarities in the Distribution of Stock Market Price changes between the Eighteenth and Twentieth Centuries. *Journal of Business*, 71(1), 55-79.
- Harrison, P. (2001). Rational Equity Valuation at the Time of the South Sea Bubble. *History of Political Economy*, 33(2), 269-281.
- Harris, M. and A. Raviv (1993). Differences of opinion make a Horse race. *Review of Financial Studies*, 6(3), 473-506.
- Harrison, J.M and D. Kreps (1978). Speculative Behavior in a Stock market with Heterogeneous Expectations. *Quarterly Journal of Economics*, 92(2), 323-336.
- Heaton, J.B. and N.G. Polson (2011). Smart Money, Dumb Money and Learning from Prices. *Working Paper, University of Chicago*.
- Hong, H., and J. C. Stein. (2003). Differences of Opinion, Short-Sales Constraints, and Market Crashes. *Review of Financial Studies*, 16, 487-525.
- Jacquier, E. (2008). Long-term forecasts of mean returns: Statistical versus Economic rationales. *Working paper, HEC Montreal*.
- Jacquier, E., Kane, A., and A. Marcus (2003). Geometric or Arithmetic Mean: A Reconsideration. *Financial Analysts Journal*, 59(6), 46-53.
- Jacquier, E., Kane, A., and A. Marcus (2005). Optimal Estimation of the Risk Premium for the Long-Term and Asset Allocation. *Journal of Financial Econometrics*, 3, 37-56.
- Jacquier, E. and N. G. Polson (2011). Bayesian Methods in Finance. *Handbook of Bayesian Econometrics*, H. van Dyk et al (eds).
- Jacquier, E., N. G. Polson and P. Rossi (1994). Bayesian Analysis of Stochastic Volatility Models. *Journal of Business and Economic Statistics*, 12(4), 371-89.
- Jacquier, E., N. G. Polson and P. Rossi (2005). Bayesian Analysis of Stochastic Volatility Models with Fat-tails and Correlated Errors. *Journal of Econometrics*, 122(1), 185-212.

- Johannes, M. and N. G. Polson (2009). MCMC Methods for Continuous-time Financial Econometrics. *Handbook of Financial Econometrics* (eds Ait-Sahalia and L. P. Hansen), 1-72.
- Johannes, M., A. Korteweg and N. G. Polson (2011). Sequential Learning, Predictive Regressions and Optimal Portfolio Allocation. *Working paper*.
- Kelly, J.R. (1956). A new interpretation of the Information rate. *Bell System Technical Journal*, 35, 917-926.
- Markowitz, H. (2006). de Finetti scoops Markowitz. *Journal of Investment Management*, 4(3), 5-18.
- Merton, R.C. (1969). Lifetime Portfolio Selection under Uncertainty: the continuous time case. *Review of Economics and Statistics*, 50, 247-257.
- Morris, S. (1996). Speculative Investor Behaviour and Learning. *Quarterly Journal of Economics*, 111(4), 1111-1133.
- Odean, T. (1998). Volume, Volatility, Price, and Profit When All Traders Are Above Average. *Journal of Finance*, 53, 1887-1934.
- Polson, N.G. and B.V. Tew (2000). Bayesian Portfolio Selection: an empirical analysis of the S&P500 indices 1970-1996. *Journal of Business and Economic Statistics*, 164-173.
- Ramsey, F.P. (1926). Truth and Probability. Reprinted in *Studies in Subjective Probability*, Kyburg and Smokler (1980) (eds).
- Savage, L.J. (1954). *Foundations of Statistics*. John Wiley and Sons.
- Scheinkman, J. A., and W. Xiong (2003). Overconfidence and Speculative Bubbles. *Journal of Political Economy*, 111(6), 1183-1220.
- Siegel, J.J. (1994). *Stocks for the Long Run*. McGraw-Hill.
- Stein, J. C. (2009). Sophisticated Investors and Market Efficiency. *Journal of Finance*, 64(4), 1517-48.
- Varian, H. R. (1985). Divergence of Opinion in Complete Markets: A Note. *Journal of Finance*, 40(1), 309-317.
- Varian, H. R. (1989). Differences of Opinion in Financial Markets. In *Financial Risk: Theory, Evidence and Implications*. Courtnet C. Stone (ed).

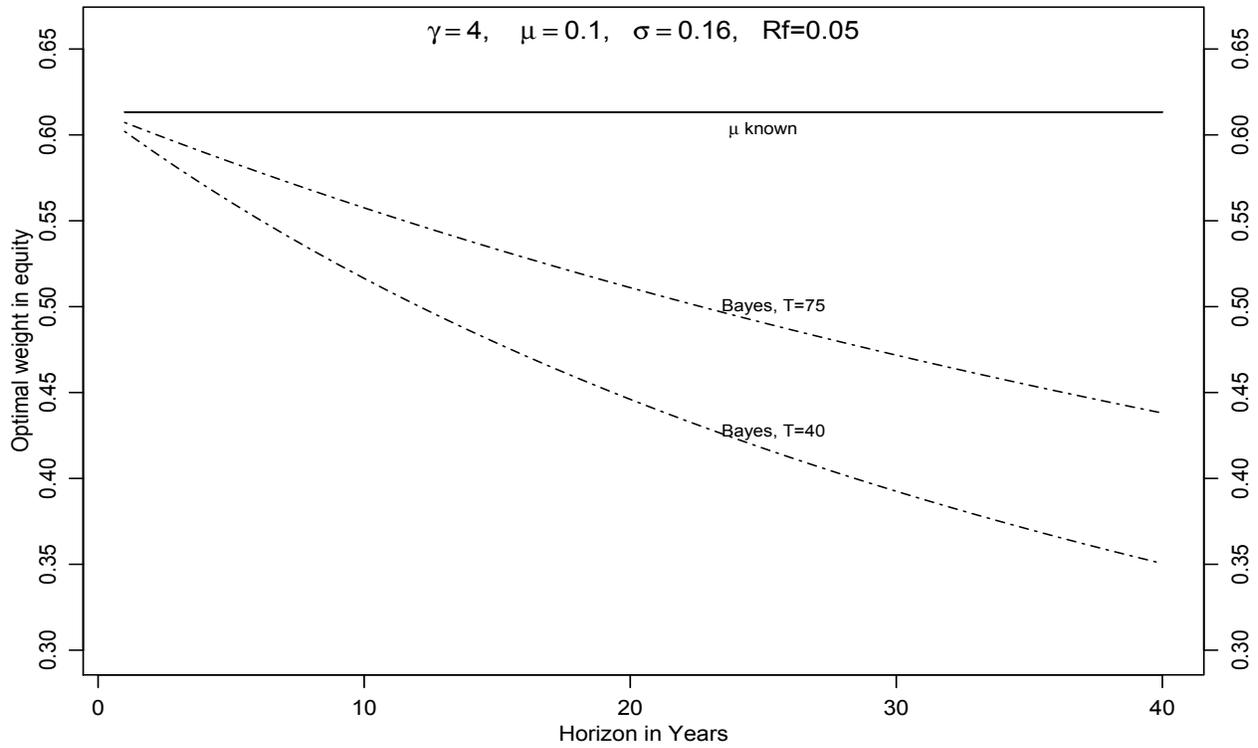


Figure 1: Bayesian Long-term asset allocation under uncertainty
Diffuse prior on μ .

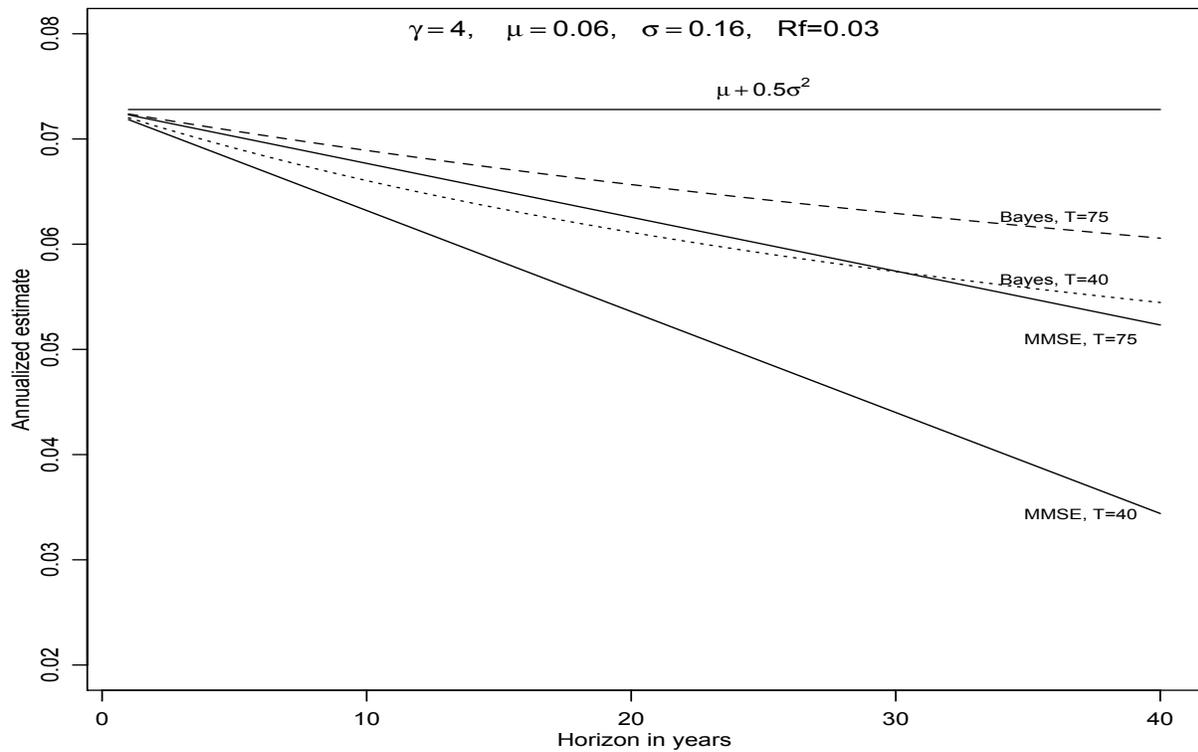


Figure 2: Bayesian and classical minimum mean-squared estimates of annual compound factor